


CONNEXIONS



The Interoperability Report

March 1990

Volume 4, No. 3

Connexions—

The Interoperability Report tracks current and emerging standards and technologies within the computer and communications industry.

In this issue:

Profile: AARNet.....	2
Book Review.....	11
Clarification of GOSIP.....	13
Automatic Configuration of Internet Hosts.....	15

Connexions is published monthly by Interop, Inc., 480 San Antonio Road, Suite 100, Mountain View, CA 94040, USA. 415-941-3399. Fax: 415-949-1779.

Copyright © 1990 by Interop, Inc.
Quotation with attribution encouraged.

Connexions—The Interoperability Report and the Connexions masthead are trademarks of Interop, Inc.

ISSN 0894-5926

From the Editor

Previous issues of Connexions have featured profiles of large internets. This month we have another such profile, only this time the network doesn't actually exist (yet). The Australian Academic and Research Network (AARNet) is currently being constructed. The article—which is by AARNet's technical manager, Geoff Huston,—outlines the objectives, network architecture, and implementation program for AARNet.

At the most recent Internet Engineering Task Force (IETF) meeting in Tallahassee, Florida, Phill Gross (IETF Chair) gave a brief presentation on the Government OSI Profile (GOSIP) as it applies to the Internet. Connexions asked for a short write-up on this topic. The article was co-authored by Phill Gross and Rebecca Nitzan and appears on pages 13 and 14.

The AARNet and GOSIP articles give me an opportunity to give you a little insight into how Connexions is produced. Articles normally arrive via e-mail to the Editor's Internet mailbox. From there the (ASCII text) files are transferred to a Macintosh computer, using the *Kermit* file transfer program. Once the text has been stored in the Macintosh it is edited and formatted for final camera-ready copy, and printed by Globe Printing Company in San Jose, California.

When I first received the AARNet article, it did not contain any illustrations. I sent a message to Geoff Huston asking how we might incorporate a network map and a picture of the network structure. Within a few minutes he pointed me to a couple of *PostScript* files which I fetched via FTP from Australia and printed locally. In the case of the GOSIP article, I had a quick question and tried calling Phill Gross, only to discover that he was on vacation in Florida, carrying his laptop. The question was quickly resolved via e-mail which he read from his hotel room, using the laptop as a terminal. The wonders of modern technology!

In the September 1989 issue of Connexions, Greg Minshall lamented the complexity of IP networks as compared to AppleTalk networks which are very much "plug-and-play." This month, Ralph Droms discusses the problem in more detail, and explains what is being done in the Internet community to develop solutions.


Reflecting the fact that *interoperability* is really what we're all about, Advanced Computing Environments will change its name to **Interop, Inc.**, effective March 1, 1990. Our address and telephone number remains the same. (See page 12).

**Profile: AARNet—
The Australian Academic and Research Network
by Geoff Huston, AARNet Network Technical Manager**

Introduction

The implementation of a National Research and Academic Network has been a matter of growing importance to the Australian higher education sector and various research institutions for some years. Over this period Australia has seen similar projects established and intensively used in peer nations, and the present lack of an Australian national network infrastructure for academic and research use is very much an anomalous situation.

Since 1982 various aspects of the establishment of such a network have been the subject of attention from the academic and research community, with the effort to establish a national network intensifying over the last two years. It has now been almost 12 months since the initial report advocating immediate support for the establishment of a national academic and research computer network within Australia was presented to the higher education sector for consideration.



After further widespread consultation the design specifications of the network have been prepared, together with an accompanying blue-print for the network's implementation. Additionally the ongoing activities of operational management, and planning and policy determination have been considered by the committees. Recommendations on the network's design and management structures have been endorsed by the relevant academic and research executive bodies in mid-1989.

Benefits

The benefits of this program to Australia's research efforts lie in the fostering of national and international collaborative research efforts through a common communications service; the direct access to information sources and other resources across the network; and the direct impact on the productivity and effectiveness of research and academic activities by providing access to national and international communications facilities to the desktops of researchers and scholars.

In doing so the program can provide a communications infrastructure within Australia which can lift this area of research support to a level comparable to peer nations, with the consequent strategic benefits to the nation that a competitive and productive research environment can provide.

Background

The last two decades have seen a remarkable evolution in computing technology, and the associated area of computer networking has been no exception to this technological development. Research activity on large scale heterogeneous computer networks commenced over two decades ago with one of the major catalysts being the United States' *Advanced Research Projects Agency* (ARPA) initial sponsorship of the ARPANET project. Today the resultant *Internet* network effectively spans the academic and research environment within the United States. Many other nations have implemented similar national networks to support research and academic activities; the United Kingdom Computer Board has been actively supporting the *Joint Academic Network* (JANET) network project for the last decade.

The *European Academic and Research Network* (EARN) serves the European community by providing direct links to member institutions as well as interconnecting a number of national networks through Europe including *INRIA* (France), *UNINETT* (Norway) and *SUNET* (Sweden).

Outside of Europe and the United States, similar national networks for research and academic purposes have been implemented in Japan, Korea, and New Zealand, and programs are underway in many other nations. There has been a parallel activity of providing international linkages between these networks, providing to each user within each network not only access to other users within the local network, but also global access across a large heterogeneous computing environment.

Overseas experience with such an infrastructural facility has pointed to productive benefits on a national scale, in terms of collaborative research efforts, drawing academic and commercial research to areas of common focus, enabling direct access to leading edge research tools and direct and immediate access to one of the crucial "raw materials" of research effort—information in all its varied forms. Indeed the view from those nations which have established such national research networks is that access to a high performance computer network is an indispensable component of the overall national infrastructure required to perform high quality research. The conclusion drawn from observing the results of these overseas programs is that the requirement for such a facility to be established within Australia is crucial to our continued participation in many aspects of leading edge research activity.

The pace at which these benefits can be realized within the academic and research community is extremely rapid: this is indicated by the rate of growth of the usage of such networking facilities in overseas research networks. As a typical example, the *NSFNET* (funded by the National Science Foundation of the United States) has been experiencing a network usage load which has doubled every 4 months throughout its three years of operation, and *NSFNET*'s data traffic capacity will have increased by more than one thousandfold by mid 1990 from the initial 1986 levels.

continued on next page

Profile: AARNet (*continued*)

Overview of activities

While the proposal for a national network has for some years been the subject of investigation and debate within the academic and research community as to the most appropriate way in which such a project should be undertaken, it is only within the last two years that a number of essential pre-conditions for such a network have been satisfied. Such pre-conditions include the extent to which individual Universities and Colleges have proceeded with computer networking within each local campus; the decreasing real cost of data communications technology; the increasing adoption of open solutions for networking services and applications on each campus; and the integration of the use of computer networks within many functions of the higher education institutions.

A report proposing the establishment of a national academic and research network was prepared for Australian higher education executive bodies in late 1988.

Following consideration of this report further activities have concentrated on refining this broad proposal and drafting of a detailed strategy as to the most appropriate way in which such a network could be established and subsequently operated and managed.

The technical aspects of such a network facility were discussed in detail at the 3rd Australian Academic and Research Network Workshop in December 1988, and the workshop proposed the adoption of an multi-protocol wide area network architecture as a viable and universally acceptable methodology for the implementation of the network.

A submission is also being made to the Australian Research Council for assistance in funding this network. The submission calls for Governmental commitment to fund the national and international components of the network as part of the Government's ongoing commitment to support academic and research activity within the nation.

Objectives

The objective of the Australian Academic and Research Network is the provision of a high speed communications network to the members of the Australian academic and research community. This will involve higher education institutions, the *Commonwealth Scientific and Industrial Research Organisation* (CSIRO) and many other research areas throughout Australia.

Activities proposed to be supported on the network include:

- *Collaborative Effort:* The objective here is to facilitate collaboration on a national scale, and also allow Australian researchers to collaborate in international research efforts. This would encompass the ability to exchange information, software and computer data between users of the network, enabling the support of geographically dispersed research groups with a common focus of activity.
- *Electronic Mail Delivery:* The ability to communicate with peer users both nationally and internationally using a fast and reliable electronic mail delivery system.
- *Information Access:* The ability to access information on remote systems, through either direct remote interactive access or through distributed database applications.

This would encompass direct access to discipline-specific information sources, access to library catalogues, and similar. This activity can also facilitate the rapid dissemination of research methods and results throughout the research community.

- *Supercomputer Access:* The ability to access remote supercomputer facilities in a productive fashion using local networking resources to access such supercomputers across the national network.

The result of such a network is the creation of a distributed computing environment where each computing system or workstation can be used within a local, national and global networking environment to access other users or remote resources which are accessible over the network, and to provide the ability to publish local resources, information, software or data for access by other network members.

Gathering together these various networking activities, the overall intent is to provide a national communications environment which can enable collaborative and productive effort on a national scale amongst the nation's researchers and scholars.

From the academic perspective the network's objectives are to construct the basic infrastructure for services and applications which can address many of the current and anticipated communications requirements of higher educational institutions. This may well include aspects of networking support for distance education programs and support for tertiary administration activities as well as the areas of peer communications and information access and retrieval.

International collaboration

Also important are the aspects of international connectivity and compatibility. It is recognised that international collaboration is playing an increasingly important role within all areas of research activity, and for this reason good international connectivity to peer research networks is a major aspect of this network's objectives. For this reason the network design includes the objective of ensuring that Australian users can directly access overseas networks using protocols and tools which provide access to overseas facilities, and overseas users can perform the reverse operation into the Australian network. In the longer term, Australia must be prepared to take an active role in setting up international network links within the local Pacific and South East Asian area. The trans-Pacific link to the United States via Hawaii is perceived to be the forerunner to potential links to New Zealand, Japan and other South East Asian nations.

Architecture

The AARNet network design methodology has been to nominate a network architectural approach to address the various issues involved in the construction and maintenance of a national network. The goals of this network architecture are:

- To use existing networking technology;
- To use the expertise existing within the network member sites;
- Be readily implementable;
- Have architectural simplicity and uniformity;
- Be compatible with existing Australian and International networks;
- Allow for evolution of technology

continued on next page

Profile: AARNet (*continued*)

These architectural goals can be achieved by the approach of using the *Local Area Network* (LAN) as the basic connection unit of the network, and constructing a *Wide Area Network* (WAN) to provide packet delivery services between these LANs. The architecture also nominates a set of supported networking protocols, and uses protocol-specific network routing support on each link to provide maximum data capacity for each connection. The LAN-based approach mirrors exactly the network approach being currently implemented by the majority of members of the academic and research community in providing network services within each site.

Network engineering

This architecture effectively places a compatible superstructure of a national network on the top of the model of such local networks, allowing those networking services already used within each site to be extended into a national wide area domain. In this way this proposal effectively utilises one of the major assets within this community which has been developed within each institution: the technical expertise in constructing uniform network services for an institution to interconnect a diverse set of computing resources. The design of the network is the result of consideration of a number of engineering factors, which are briefly outlined here:

- *Availability of cost-effective communications technology.* The network will be built using commercially available equipment.
- *Availability of transmission bandwidth for digital communications.* The communications links are based on the availability of terrestrial digital links. These links are available as either low speed (9.6K), mid speed (48K–64K) or high speed (2M) links, and the communications tariffs are based on a function of the bandwidth and the distance of each link. Consideration has been given to using the public X.25 packet switched network (Austpac), but the problems of interfacing connectionless traffic into an incrementally charged connection-oriented carrier network preclude this option on the grounds of operational cost and technical issues.
- *Projections of network usage levels.* And the matching of the network capacity to the demands placed by usage.
- *Network reliability.* The network should avoid the inclusion of critical points of failure. In order to provide a reliable service to users alternate paths need to be configured on critical routes to minimise the disruption caused by line or equipment failure.
- *Network performance.* The performance of the network is related to the traffic usage of the network, the communications bandwidth of the links and the link termination equipment, and the size of the network.
- *Network management.* The engineering goal here is to be able to maintain a complete picture of the operational status of the network from a number of monitoring points within the network. Also required is the ability to monitor the usage profile of the network, logging information regarding the utilisation of each network component both in terms of actual load levels and also user application type. Such information is a critical component of the management information base of the network.

- *Network Extensibility.* A major engineering consideration is to modularize the network architecture such that the issues of the physical communications links can be viewed independently from the superstructure of the network routing and application technologies, allowing variations in the type and bandwidth of digital links without the requirement to alter the remainder of the network infrastructure. This modularity also is applicable to the nomination of supported network protocols, and the network should allow the introduction of support for additional protocol stacks in a modular fashion. The overall intent is to minimize the costs involved in ensuring that the network can provide an ongoing commitment to meet the requirements of the user community on both capacity and supported services.

Implementation program

The objective of the network implementation program is to rapidly establish a high performance national computer network which provides a set of common communications services to scholars and researchers throughout the nation.

This high performance network will comprise a common backbone network and eight State Regional networks. The backbone network is designed to carry data traffic between each Regional network and also provide the interface between Australia and peer international networks. The eight Regional Networks are configured with interfaces from each State Capital city (Brisbane, Sydney, Canberra, Melbourne, Adelaide, Perth, Hobart, Darwin) into the backbone network. The backbone network is designed to support the following network usage estimates and performance parameters:

- *Distributed access* from all states to major high performance computing platforms hosted in sites in Melbourne, Sydney, Canberra and potentially Brisbane. Such access includes remote interactive access from terminals and workstations, remote job submission, and large data set transfers.
- *Remote interactive access* between any two points on the network with minimal transmission delays imposed by the network.
- *Efficient transfer of electronic mail and news* throughout the network.
- *Background file transfer.*

Links

Taking into account estimated traffic levels for these activities, and including internal traffic overheads (for facilities such as name servers, etc), there is an immediate requirement for high capacity trunk lines (2 Mbps) to service the South East sector of Australia, and within three years this requirement for capacity will extend to Queensland, South Australia and Western Australia. The remaining trunk lines are estimated to require aggregate capacity of between 144Kbps to 1Mbps over the same period.

Within each Regional network the bandwidth of the regional links will vary greatly. It is anticipated that within this same three year period the larger higher education institutions will require 1–2 Mbps aggregate bandwidth, while the smaller institutions will use regional links of bandwidth between 9.6K bps to 144Kbps.

Profile: AARNet (*continued*)

The major international link requirement is a connection into the United States research networks (which in turn have high capacity connectivity through to the European academic and research networks). This implementation program includes the upgrading of the bandwidth of the current 64Kbps cable link from Melbourne to Hawaii to 128Kbps bandwidth, or possibly greater, depending on recorded traffic trends.

The proposed implementation strategy to achieve this network configuration is an approach of phased introduction of high capacity network facilities over the next three years.

It is recognised that there is a requirement for an extensive effort on the part of all member sites in order to configure their local systems to be a peer member of the national network. Member sites also need to inform all users of the ways in which the local facility may be productively utilized in terms of the network services and applications as well as the methods of access to remote network resources. Therefore it is anticipated that it will take some months following the initial installation of network connectivity for the usage rates to reach production levels. Accordingly the physical network implementation program includes this phased introduction of bandwidth and capacity to keep pace with increasing usage of the network. Additionally this approach will allow the network to effectively utilize the forthcoming high performance digital communications technologies being introduced by the national carriers over the next few years.

Network development

The initial phase of the network will establish the basic trunk national and regional links using star topologies to provide the interconnectivity of the Regional Networks, and to provide institutional interconnectivity within each Regional network.

For the National Backbone network the initial configuration includes high capacity trunk lines to link the Regional networks of New South Wales, the Australian Capital Territory and Victoria. Other Regional networks will be served in the first instance by mid-speed links. (See map on page 2).

Subsequent phases of the National Backbone will extend this high capacity service to the Regional networks of Queensland, South Australia, Western Australia and (using multiple mid-speed links) Tasmania. Additionally these subsequent phases will provide additional circuits for enhanced overall reliability of the backbone network by removing the absolute reliance on a single National network hub to provide Regional connectivity.


For the Regional networks the initial configuration uses mid-speed lines radiating out from a Regional hub to interface units located within each institution's local area network (these networks will also utilize existing high capacity links where already installed). The communications equipment to be used within these regional networks is to be identical to that employed on the National network, ensuring that the network services provided within each Regional network are identical to those available at a national level.

Subsequent phases of the Regional networks will include the installation of additional capacity into the networks.

The exact program within each phase will be determined by the measured usage levels of the network links and the available funding for each phase.

Protocol support

The network design includes the support for a number of different network protocols to coexist within a single infrastructure of physical communications links. The initial phase of the network will support three protocol stacks; ISO OSI (using CLNS Layer 3), TCP/IP and DECnet Phase IV. Attention will also be given to the appropriate mechanisms to support access into the international Packet Switched Networks using the X.25 interface protocol once the initial phase of the network has been set into production.



In the longer term this supported protocol set may be reduced: with Digital Equipment Corporation's announced migration of DECnet Phase V to use the ISO OSI stack for the lower 4 protocol layers (with CLNS Layer 3), and the intended migration strategy in the U.S. of the TCP/IP protocols into ISO OSI (again with CLNS Layer 3). The long term objective of migration to an ISO OSI compliant network structure is achievable from this starting position.

Network services

The services provided by the network are effectively an extension of most of those services already available on each site's LAN. With the increasing use of a distributed model of computing support within the academic and research environment, focusing on workstations and network access to common resources, there has been a rapid growth in the use of network applications within each site, and a corresponding increase in the expertise in the productive use of network applications within the user population. Such LAN network applications allow users to exchange software and data, access remote systems, submit jobs for remote execution, and link network resources such as printers, disks, and processors to the local host system.

continued on next page

Profile: AARNet (*continued*)

The applications available over the WAN are no different in many respects—the limitations of the WAN are essentially those of bandwidth, so that some applications, such as disk sharing, are not generally viable over a WAN. However the majority of network applications in local use on a LAN will run unaltered in the WAN/LAN combined environment.

The following is a brief (and by no means complete) list of some of the applications which can be supported over this WAN/LAN network, and the ways in which they are commonly used:

- *Electronic Mail.* The interchange of ideas, information and resources between users who are linked by a common electronic mail system is perhaps one of the most well known of network services for many users. Within Australia many users are already linked by a number of mail services, including ACSnet, PSImail, Coloured Book Mail and such. The network will allow faster delivery of mail by allowing systems to directly exchange mail messages without the delays of intermediary store/forward mail transfer agents, and with the introduction of mail gateways, will allow each local mail interface the ability to directly address mail to any other user on the network.
- *Remote Access* is the other highly visible application of networks; supporting the ability to allow a local user to establish an interactive terminal session on a remote host. This facility is used for a wide variety of purposes, including remote access to catalogues, databases and other information sources, as well as the ability to access a shared facility from a remote location.
- *News.* As well as mail transfer, the other major component of messaging networks is the exchange of public notes between users on a network of host systems, creating a resultant network-wide bulletin board. The *USENET* news network in the United States currently delivers 3,000 new messages per day to a global readership estimated to be of the order of 1,000,000 readers. NEWS is used to distribute software and software updates, provide technical assistance on a peer basis on a wide variety of subjects, and to allow the interchange of ideas on a wide range of subjects of particular relevance to the research and academic community. Within Australia the ACSnet network delivers less than a quarter of these messages to ACSnet sites. The AARNet has the network capacity to deliver the complete USENET set throughout Australia, and also improve the news delivery times.
- *File Transfer.* The ability to transfer computer files through a network is also an area of common usage. Such file transfer facilities enable the sharing of resources and information between users, and also provide a mechanism for the rapid distribution of software, applications and data.

There are many other network applications not mentioned above, including remote job submission, directory services, electronic document exchange, distributed databases, distributed filing systems and such.

As with the trend within local sites towards a distributed computing model as a more productive and cost effective computing strategy than a single central computing resource, the academic and research environments are now in a position to take advantage of the significant opportunities to productively utilise a national and global distributed computing environment.

**Planned activities
for the AARNet**

The initial implementation phase is expected to be completed by April 1990, providing connectivity based on the Internet protocol suite. A nation-wide DECnet will be implemented by June 1990 using the Phase IV routing functionality of DECnet, and OSI routing functionality will be added in the last half of 1990 using CLNS services. Additional user-level services will also be added over this period, including planned access into the Australian Bibliographic Network. In the second phase, scheduled for 1991, additional 2Mbps trunk links will provide additional bandwidth of the Backbone network, complemented by selective upgrading of the tail end links into a number of larger sites.

GEOFF HUSTON received his B.Sc (1978) and M.Sc (1983) from the Australian National University (ANU). For the past ten years he has worked in for the Computer Services Centre of the ANU in both systems support and systems management roles. He authored the VAX/VMS implementation of the USENET NEWS application and is still providing support for this software over the networks. Since March of 1989 he has been employed by AARNet as the Network Technical Manager, and has been largely responsible for the planning effort associated with bringing this network into existence.

Book Review

Standards for Open Systems Interconnection, by Knowles, Larmouth and Knightson. First published by BSP Professional Books in 1987, ISBN 0-632-01868-2, 388 pages.

Audience

The foreword in the book contains the best description of the intended audience:

- What is OSI?
- Why might I wish to use OSI?
- How do I go about implementing OSI?
- How should I approach procurement of OSI implementations?

Organisation

This book is now three years old. Do not let this put you off. Despite the fact that standards have a habit of never lapsing into traditions, the authors were sufficiently well informed that much of the material is still accurate. The book is divided into five parts:

An overview and background section covers the standard material on the Organisations and Processes involved in standards, and the all pervasive Reference Model. This includes several illuminating pages on the Connectionless Addendum and the Naming and Addressing Addendum. If you were ever confused the use of terms such as "relay," or "title," "SAP," and "Selector," then this is well worth a read.

continued on next page

Book Review (*continued*)

The second part covers Interconnection of Open Systems in two chapters, layers 1–3 then layer 4. The use of State Transition Diagrams and Time Sequence pictures make the actual protocol operations fairly clear. These chapters include packet formats and the infamous OSI Network Address Format.

The third part covers Common Services. In three chapters this covers Session (the best justification you will find for Session aside from Jam), Presentation (the often forgotten protocol as well as the Abstract Syntax Notation 1), and Application Service Elements (well, Association Control and Commitment, Concurrency and Recovery).

The fourth part is about applications. Unfortunately, this only covers FTAM, JTM and VTP, omitting the rather more interesting X.400/MHS/MOTIS. Since FTAM is superseded by de facto distributed file systems such as NFS/RFS/Andrew, JTM is not widespread and VTP is overcome by X-Windows, this is the least useful section. However, it is well presented and some of the ideas in JTM will be useful to those working in Distributed Execution Environments.

The last part is called "Wider Issues," which means (as usual) that it covers all the things that are absolute requirements in real networks but only standardised at the end of the day: Management, Security and Conformance.

Management and Security

The sections on Management and security are short, and better dealt with elsewhere: M. Sloman (Ed.) *Management of Local Area Networks Report*, Part 2 of COST11 Bis Report on LANs, October 1984, and D. E. Denning, & P. J. Denning *Cryptography and Data Security Institute*, Addison-Wesley, 1983, ISBN 0-201-10150-5.

However, the section on Conformance, Conformance Testing and Procurement is useful (though brief) and is the only explanation of PICS (Protocol Implementation Conformance Statements) that I have found in the literature.

Conclusion

I find this a useful book as a quick reference to standards. It is reasonably thorough without being verbose, and is accurate. It would be a valuable alternative to such books as Halsell's *Data Communications, Computer Networks and OSI*. It is certainly not a book for the implementor (despite the foreword), rather for the manager who would like to be aware of what all the acronyms stand for and the principles (and principals) behind OSI. —Jon Crowcroft

Interop, Inc.
480 San Antonio Road
Suite 100
Mountain View
California 94040, USA

How to reach Interop, Inc. by telephone:

- In the US, toll free: 1-800-INTEROP
- In England, toll free: 0800-891-464
- In France, toll free: 19-0590-1132
- In Germany, toll free: 0130-81-1296
- In Italy, toll free: 1678-74-006
- Elsewhere in the world: +1 415-941-3399
- By FAX: +1 415-949-1779

Clarification of GOSIP

by Phill Gross (NRI) and Rebecca Nitzan (NASA)

Background

The *Government OSI Profile* (GOSIP), issued as FIPS 146 by the National Institute of Standards and Technology (NIST), specifies the details of OSI for use in the U.S. Government.

OSI “profiles” are important because OSI standards allow many potential *options* and *choices*. Without careful specification and prior agreement, different vendor products might very well conform to the OSI standards but not interoperate with each other. Therefore, a major goal of FIPS 146 is to insure that the U.S. Government be able to buy interoperable OSI products from different vendors.

The first version of GOSIP was published in August 1988 following a comment period beginning in early 1987. GOSIP was adopted as FIPS 146 in February 1989 and will become a Federal procurement requirement in August 1990 [1]. GOSIP was written by an inter-agency group and continues to evolve under the guidance of the GOSIP Advanced Requirements Group. A second version of GOSIP will become a FIPS in the summer of 1990 and will then become a Federal procurement requirement 18 months later [2].

GOSIP Users' Guide

There is an additional publication called the *GOSIP Users' Guide* which provides an expanded explanation of GOSIP including tutorials, interpretation, integration planning advice, and information on registration [3]. The GOSIP Users' Guide will be updated and re-released in coordination with each version of GOSIP.

The Internet Activities Board (IAB) and the Internet Engineering Task Force (IETF) are fully committed to integrate OSI into the Internet. In particular, one of the eight technical areas of concentration in the IETF is devoted to OSI integration, and the IETF is represented on the GOSIP Advanced Requirements Group.

Source of confusion?

GOSIP is an important tool for planning OSI integration. However, as the August 1990 requirement date for GOSIP compliance approaches, there has also been an increasing amount of concern as to how GOSIP should be applied to near-term network planning. In particular, there appears to be a common misunderstanding that GOSIP mandates a transition to OSI beginning in August 1990.

For example, in the January 1990 *IEEE Spectrum* (Technology '90), there is the following quote in the section on Data Communications (page 35-36):

“OSI protocols are viewed as a long-term answer to the problem. But, the scarcity of products on the market hinders devising a network strategy around OSI...”

Many vendors are still pinning their hopes for OSI on FIPS 146 (GOSIP), which requires that Federal agencies start using OSI products after August.”

GOSIP does not “require” that Federal agencies start using OSI products after August 1990. GOSIP is a procurement specification. GOSIP does not mandate, or even explicitly address, the issue of protocol transition.

continued on next page

Some clarifying points about GOSIP

Summary

As a procurement specification, GOSIP does not apply to existing installed equipment. It applies to new network procurements and major upgrades to existing networks. "Major upgrade" does not necessarily apply to increasing the number of components in existing non-GOSIP networks. When GOSIP does apply, it is not exclusionary. That is, other protocol families can continue to be procured and used. When GOSIP does apply, waivers are allowed in consideration of specific agency requirements. When GOSIP does not apply, no waiver may be necessary.

Agencies have the responsibility for developing their own waiver process, and for determining the applicability of GOSIP to any specific procurement. NIST does not have an enforcement role regarding GOSIP. In general, agencies are responsible for developing their own agency-wide plans for GOSIP compliance in their network procurements.

References

The large existing installed base of TCP/IP and other protocol users, the limited availability of commercial OSI products, and the still incomplete development of OSI standards (e.g., for routing, network management, and directory services) combine to make a near-term transition to a ubiquitous OSI environment in the Internet unrealistic. GOSIP is an important tool for procuring interoperable OSI products for the U.S. Government. However, GOSIP does not mandate, or even explicitly address, the issue of OSI transition.

The points in this article will be given in more detail in a forthcoming RFC [4] issued by Vint Cerf (IAB) and Kevin Mills (NIST).

- [1] "U.S. Government Open Systems Interconnection Profile," August 1988, U.S. Federal Information Processing Standards Publication (FIPS) 146, Version 1.
- [2] "U.S. Government Open Systems Interconnection Profile," April 1989, U.S. FIPS 146-1, Draft Version 2.
- [3] Tim Boland, "Government Open Systems Interconnection Profile Users' Guide," August 1989, NIST Special Publication 500-163.
- [4] "Explaining the Role of GOSIP," RFC (xxxx), V. Cerf & K. Mills.

PHILL GROSS joined the Corporation for National Research Initiatives in November 1988. In 1986, he was a founding member of the Internet Engineering Task Force (IETF) of the Internet Activities Board (IAB). He has served as IETF Chair and IAB member since July 1987. Since joining NRI, Mr. Gross has organized and now chairs the FRICC Engineering Planning Group, which is responsible for engineering joint-agency and international network connections. Prior to joining NRI, he was at the MITRE Corporation where he developed a joint-agency program in networking performance research. Mr. Gross received his masters degree in computer science from the Pennsylvania State University.

REBECCA NITZAN has been working for Sterling Software since September 1989 on OSI implementation issues for the NASA Science Internet (east coast) Project Office. She co-chairs the FRICC OSI Planning Group, which coordinates FRICC agency OSI integration and is a member of the IETF. Prior to joining the NASA project, Rebecca was at Lawrence Livermore National Laboratory working on the implementation of DOE's Energy Science Network (ESNET), which included the design and implementation of router and network management software. She received her B.S. degree in Computer Science and in Mathematics from the University of the Pacific, California in 1985.

Automatic Configuration of Internet Hosts

by Ralph E. Droms, Bucknell University

The problem

The exponential growth of the Internet has greatly increased the availability of TCP/IP based network communications. Along with that access to network communications comes the problem of managing the hosts, routers and network hardware that make up the Internet. System managers must devise a local network architecture, interconnect physical networks with routers (traditionally known as "gateways" in the Internet community), assign identification numbers and update configuration information. Changes in the local network, such as the addition of new components or the movement of existing components to new networks, require manual revision of the network configuration information in *every* host.

The goal

Network management need not always be complicated. Apple's LocalTalk/AppleTalk protocol suite is perhaps the best known example of a "plug-and-play" network. As much as possible, AppleTalk participants require no manual intervention—the user simply connects the hardware to the network and all configuration information is obtained automatically through the network. [19]

The corresponding vision in the Internet world is for the traveling user to walk into a hotel room and connect a laptop computer to the local network through the Ethernet transceiver socket next to the phone. After some automatic configuration is complete, the laptop, `mobile1@gsu.edu`, is a full Internet participant, with an IP address, registered domain name and access to Internet services.

Why can't the TCP/IP protocol suite operate as simply as AppleTalk? There are many reasons: TCP/IP was developed in a small, static environment where the participants were primarily mainframe computers whose configuration could be managed manually and saved locally. The TCP/IP Internet is a very large, mixed environment that operates through the cooperation of many autonomous, independently managed networks. Some degree of central control is required in the Internet to provide stability and reliability. Because the Internet accommodates multi-vendor hardware from many different manufacturers, reaching agreement on a single configuration mechanism is difficult.

Details

It's useful to talk about Internet participants as hosts and routers. A *host* is an end participant in an Internet communication, while a *router* interconnects physical networks. Hosts and routers each have different configuration requirements.

Hosts require configuration information for several different network functions. At the hardware layer, a host must know the physical address of any other local systems with which it exchanges frames. To participate in Internet communication, that is, to exchange IP packets with other Internet participants, a host must know the local network number (and several related parameters), the host's own IP address on the network, and the address of a router through which packets can be forwarded to other networks. Name resolution through the Domain Name System (DNS) [10, 11] requires that the host register its name with the DNS and that the host know the address of a DNS server. Finally, use of other application-level services requires that the host know the address and protocol of the particular service.

Auto Configuration of Internet Hosts (*continued*)

Routers must also have access to physical addresses, local network numbers and the router's own IP addresses (a router has a separate IP address for each network to which it is attached). In addition, a router must have topology information describing which networks are attached to which of the router's physical interfaces, and a router must have a list of other routers with which packet routing information is to be exchanged. Details of the initialization information for hosts can be found in the Host Requirements documents, RFC 1122 and RFC 1123 [3, 4] and for routers in the Gateway Requirements document, RFC 1009 [2].

Existing mechanisms

Several protocols within the TCP/IP protocol suite deal with automatic configuration. ARP [13] dynamically provides hardware addresses for local frame transmission. RARP [8] explicitly addresses the problem of IP address discovery, and allows for a dynamic IP address assignment mechanism through the extensions defined in the DRARP draft RFC [5]. ICMP [14] provides for informing hosts of additional gateways via *ICMP redirect* messages. ICMP also can provide subnet mask information through the *ICMP mask request* message and other information through the (obsolete) *ICMP information request* message.

BOOTP [6] is a transport mechanism for static configuration information. BOOTP is also extensible, and official extensions [15, 16] have been defined for much of the configuration data required by a host. Morgan proposes extensions to BOOTP for dynamic IP address assignment [12]. NIP, used by the Athena project at MIT, is a negotiation mechanism between host on a network and an assignment algorithm for dynamic IP address assignment among those hosts [17]. RLP [1] provides for location of higher level services. Sun's diskless workstations use a boot procedure that employs RARP, TFTP [18] and an RPC mechanism called bootparams to deliver configuration information and operating system code to a diskless workstation. Some Sun networks also use DRARP and an auto-installation mechanism to automate the configuration of new workstations in an existing network.

What's missing?

With all of these protocols in place, why *can't* an Internet host automatically configure itself when plugged into a network? Three working groups of the IETF are all working on parts of the problem. The *Dynamic Host Configuration Working Group* has produced a draft RFC [7] describing the host configuration problem and is developing a new protocol that will perform dynamic IP address allocation and network layer host configuration. The *Gateway Discovery Working Group* of the IETF is investigating the problem of informing hosts of local gateways and the *MTU Discovery Working Group* is developing a mechanism through which a host can dynamically determine the MTU (*Maximum Transmission Unit*) of an internetwork connection to avoid packet fragmentation.

Let's look at some specific problem areas under consideration by the Dynamic Host Configuration Working Group. First, we need to establish some ground rules:

- Make sure any new protocols are a mechanism not a policy. In particular, provide local system administrator's control over configuration where desired.
- Avoid hand configuration of hosts or of local network databases.

- Do not require a server on each physical network.
- Allow for static configuration to accommodate systems where the need for a known configuration outweighs the advantage of dynamic configuration.
- Must be upward compatible with older hosts that do not implement any dynamic configuration mechanisms, but should not be rendered unusable by the presence of dynamic configuration mechanisms on the network.

The most important problems in configuring an Internet host are getting the initial Internet layer information to the host, getting information about the Domain Name System and finding out about other high-level services. The first problem is important because it represents a “chicken-and-egg” question. If a host can only use the network to get its configuration information, but requires the configuration *before* that communication takes place, how can the host use the network for dynamic configuration? The second problem is important because access to the Domain Name System is required for resolution of mnemonic host names into IP addresses. The third problem is important because access to high-level services is what the Internet is really all about.

IP layer configuration

Of the protocols mentioned earlier, RARP is the most widely used mechanism for getting an IP address to a host. RARP does not, however, support dynamic allocation of IP addresses, so system administrators must manually enter new hosts to the RARP mechanism. RARP also does not transmit other configuration information such as the subnet mask, and requires a server on each physical network.

BOOTP is less widely used than RARP, and also does not support dynamic address allocation. The BOOTP specification does include a mechanism for forwarding requests through cooperating routers, so that BOOTP does not require a server on each physical network (in contrast to RARP). Recently proposed extensions to RARP and BOOTP address the dynamic address allocation problem.

NIP, currently in use only at MIT, provides dynamic IP address allocation, IP address transmission and subnet mask information. However, NIP, like RARP, is a link-level protocol and requires that the configuration information be provided by a source on the same physical network as the requesting host.

Where does diskless workstation initialization fit into this discussion? Users certainly don't have to enter configuration parameters every time they restart the Sun on their desk. Isn't that dynamic configuration? Well, in fact, it's not. That host re-initialization still takes a significant amount of manual intervention whenever the network structure changes; for example, when a new workstation is added or an existing workstation is moved to a new network. Remember that one of the goals of dynamic host configuration is to avoid any manual configuration of either hosts or network databases. The configuration information for diskless workstations is maintained in static files (typically, several different tables that must be hand configured) and passed to the workstation through protocols such as RARP, TFTP, bootparams and NFS.

Domain Name System configuration

There are three components to the problem of integrating a new host into the DNS. First, the new host must discover the address of a DNS server from which the host can obtain name resolution services.

continued on next page

Auto Configuration of Internet Hosts (*continued*)

Next, the new host must choose a name for itself that doesn't conflict with names chosen by other hosts in the DNS. Finally, the new host's name and IP address must be registered with the DNS.

At present, hosts usually obtain the address of a Domain Name System server through static configuration information. There is currently no mechanism within the DNS for a host to request that its name and IP address be added to the DNS database. Modifications to the host configuration mechanisms and to the DNS must be made before host can automatically make use of DNS Services.

Access to higher-level services

To use a higher level service such as the WHOIS [9] directory service, a line printer spooler or NFS, the host must find out the address of the service and the protocol the service uses. There are two components to a service's address: the IP address of the host providing the service, and the port number on that host. A host can use the Domain Name System to translate a host name to an IP address, but the host must know the name of the host on which the service is located through some mechanism. For example, the `whois` command in UNIX is written to use the host `nic.ddn.mil` by default, or the user can select a specific host.

There are two ways in which service port numbers are determined. Some port numbers, called "well-known ports" are statically assigned as Internet standards. Hosts using these services determine the port number by simply reading the number from a static table. Other port numbers are selected dynamically, and these assignments must be transmitted through some external mechanism.

There is no general directory or "yellow-pages" service that provides the names of hosts and service port numbers. Some hosts advertise their services through RFCs or by word of mouth—this is another area in which effort is required to develop automatic configuration mechanisms.

Summary

An Internet host requires a significant amount of manual configuration because of the complexity and heterogeneity of the Internet environment, and because several key mechanisms for automatic configuration are missing. In contrast, networks like AppleTalk, which are homogeneous, small in scale and support only a relatively few services, can allow automatic configuration.

In fact, the situation is more complex than the previous discussion about Internet configuration might indicate. We have begged the question about network configuration altogether—the Internet is large enough that network auto-configuration is a much more difficult problem than host configuration: security, propagation of new configuration information, variability of subnet configuration are a few of the problems encountered in dynamic network configuration. The IETF is working to solve the problem of host configuration first, and will apply the lessons learned to the broader problem of network configuration.

Some protocols and standards are already in place to provide some degree of static host configuration. Several IETF working groups are developing solutions to specific dynamic initialization problems in the IP protocol.

Much work remains in the higher layers of the Internet Protocol Suite before the Internet becomes a truly “plug-and-play” network. Interested parties can take part in this effort by joining the IETF mailing list with a request to ietf-request@isi.edu.

References

- [1] M. Acetta, “Resource Location Protocol,” RFC 887.
- [2] R. Braden & J. Postel, “Requirements for Internet Gateways,” RFC 1009.
- [3] R. Braden (Ed.), “Requirements for Internet Hosts—Communication Layers,” RFC 1122.
- [4] R. Braden (Ed.), “Requirements for Internet Hosts—Application and Support,” RFC 1123.
- [5] D. Brownell, “Dynamic Reverse Address Resolution Protocol (DRARP),” RFC Draft.
- [6] W. Croft & J. Gilmore, “Bootstrap Protocol (BOOTP),” RFC 951.
- [7] R. Droms, “Dynamic Configuration of Internet Hosts,” RFC Draft.
- [8] R. Finlayson, T. Mann, J. Mogul, & M. Theimer, “A Reverse Address Resolution Protocol,” RFC 903.
- [9] K. Harrenstien, M. Stahl, and E. Feinler, “NICNAME/WHOIS,” RFC 954.
- [10] P. Mockapetris, “Domain Names—Concepts and Facilities,” RFC 1034.
- [11] P. Mockapetris, “Domain Names—Implementation and Specification,” RFC 1035.
- [12] R. L. Morgan, “Dynamic IP Address Assignment for Ethernet Attached Hosts,” RFC Draft.
- [13] D. C. Plummer, “An Ethernet Address Resolution Protocol,” RFC 826.
- [14] J. Postel, “Internet Control Message Protocol,” RFC 792.
- [15] P. Prindeville, “BOOTP Vendor Information Extensions,” RFC 1048.
- [16] J. Reynolds. “BOOTP Vendor Information Extensions,” RFC 1084.
- [17] J. Schiller & M. Rosenstein, “A Protocol for the Dynamic Assignment of IP Addresses for use on an Ethernet,” Athena Project.
- [18] K. Sollins, “The TFTP Protocol (Revision 2),” RFC 783.
- [19] G. Minshall, “AppleTalk versus IP,” *ConneXions*, Vol. 3, No. 9, September 1989.

RALPH DROMS is an assistant professor in the Computer Science Department at Bucknell University. He is an active researcher in the area of file naming and remote file access. He is an editor of *The Journal of Internetworking*, chair of the IETF Dynamic Host Configuration working group and a member of the Internet Research Naming Group. He holds a Ph.D. in Computer Science from Purdue University.

CONNEXIONS
 480 San Antonio Road
 Suite 100
 Mountain View, CA 94040
 415-941-3399
 FAX: 415-949-1779

FIRST CLASS MAIL
 U.S. POSTAGE
 PAID
 SAN JOSE, CA
 PERMIT NO. 1

CONNEXIONS

PUBLISHER Daniel C. Lynch

EDITOR Ole J. Jacobsen

EDITORIAL ADVISORY BOARD Dr. Vinton G. Cerf, Vice President, National Research Initiatives.

Dr. David D. Clark, The Internet Architect, Massachusetts Institute of Technology.

Dr. David L. Mills, NSFnet Technical Advisor; Professor, University of Delaware.

Dr. Jonathan B. Postel, Assistant Internet Architect, Internet Activities Board; Division Director, University of Southern California Information Sciences Institute.

Subscribe to CONNEXIONS

U.S./Canada	\$125. for 12 issues/year	\$225. for 24 issues/two years	\$300. for 36 issues/three years
--------------------	---------------------------	--------------------------------	----------------------------------

International	\$ 50. additional per year (Please apply to all of the above.)		
----------------------	---	--	--

Name _____ Title _____

Company _____

Address _____

City _____ State _____ Zip _____

Country _____ Telephone () _____

Check enclosed (in U.S. dollars made payable to **CONNEXIONS**).

Charge my Visa MasterCard Am Ex Card # _____ Exp. Date _____

Signature _____

Please return this application with payment to: **CONNEXIONS**

Back issues available upon request \$15./each
 Volume discounts available upon request

480 San Antonio Road Suite 100
 Mountain View, CA 94040
 415-941-3399 FAX: 415-949-1779

CONNEXIONS